Semiannual Report: July 1992 - December 1992

Asynchronous Design for Parallel Processing Architectures

Contract No.: N00014-89-J-3036
Principal Investigator: Teresa H.-Y. Meng
CIS 132, Stanford University, Stanford, CA. 94305
Phone Number: (415) 725-3636
E-Mail Address: meng@tilden.stanford.edu

The objective of this research is to provide a design methodology for connecting heterogeneous hardware modules that have inherently different functional and timing behavior. With the constraints dictated by the system-level interaction, we need to adopt a modular design approach without compromising the global performance. The main task of this effort will be the development of the theory for optimal interface circuit synthesis from a high-level specification, with emphasis on testability and performance.

Automated Gate-Level Synthesis of Speed-Independent Circuits

As reported in our last semi-annual report, we have developed and automated a synthesis procedure to synthesize asynchronous circuits from behavioral specifications suitable for a standard-cell or gate-array implementation. The generated circuits are speed-independent; that is, they are guaranteed to be hazard-free under all possible relative delays among the gates [1]. We are currently extending the synthesis procedure to accommodate cell libraries of gates with limited fanin.

The synthesis program currently assumes a cell library which contains AND, OR, and C-element gates with unlimited fanin and inverted inputs when desired. For most specifications encountered so far these assumptions seem acceptable because the synthesized circuits are easily implementable using most standard-cell packages. In a few specifications, however, the synthesis results call for the use of AND/OR gates with more than 5 inputs in the implementation. Since most standard-cell packages don't include gates with such high-fanin, these gates must be broken up into smaller ones. Unlike synchronous circuit synthesis, the breaking-up of these high-fanin gates must be designed carefully to ensure that the resulting circuit remains hazard-free. For example, we have shown that implementing a high-fanin AND gate with an AND-gate tree produces hazardous circuits.

Breaking up high-fanin gates is a common problem to all speed-independent synthesis procedures and remains an unsolved problem in all existing comparable programs. Since larger and larger asynchronous specifications are being developed, the need to devise efficient means of breaking up high-fanin gates is a major issue in designing practical asynchronous circuits.

We have derived necessary and sufficient conditions to determine if a particular decomposition of a high-fanin gate should maintain hazard-free under the speed independence assumption. Based on these conditions, we have developed an algorithm which successfully breaks up high-fanin gates without introducing possible circuit hazards. The algorithm is not complete, however, in that we haven't found *the necessary and sufficient* condition on a circuit's behavioral specification to ensure the existence of a gate-level implementation in which each gate is of limited fanin. The derivation of this condition and means of achieving it will be the main focus of our work in the coming months. As a side effect of synthesizing correct circuits, we have also constrained the algorithm to optimize for circuit area.

Synthesis of Timed Asynchronous Circuits

Timed asynchronous circuits have recently gained much attention because of the increasing need for mixed synchronous/asynchronous designs. As of the last semi-annual report, we had developed a methodology that utilizes timing constraints to synthesize timed asynchronous circuits. We begin with a cyclic graph specification to which timing constraints have been added. The cyclic graph is unfolded into an infinite acyclic graph. We show that by analyzing two finite subgraphs of the infinite acyclic graph we can efficiently detect and remove redundancy from the specification [2]. With the reduced specification, we systematicly synthesize a minimal-hardware implementation based on the given timing constraints.

In several practical examples, we have shown that significant reductions in circuit complexity can be achieved using conservative timing constraints. In particular, in a memory management unit, we were able to reduce the circuit complexity by over 50 percent over the speed-independent implementation [2].

Recently, our emphasis has been focused on enlarging the class of circuit specifications that can be synthesized. Our timing analysis algorithm had assumed that that no event occurs more than once per cycle. This led to a restriction of only one rising and falling transition of a signal per cycle. To remove this restriction, we give each occurrence of an event in a cycle a unique name. During timing analysis, these signals are treated separately, but during synthesis they are recombined. Based on this simple technique, we have been able to synthesize circuits of multiple transitions [3].

On the issue of non-deterministic circuit specifications, our original timing analysis had assumed that no such behavior is allowed. We have relaxed this restriction, and can now synthesize some examples with a limited degree of input choice. This is accomplished by transforming a specification with input choice to a single cycle without choice but exercising each possible choice sequentially. Using this approach, we are able to synthesize a reasonably complicated DRAM controller. This example is also interesting in that it demonstrates how our approach can be used to synthesize mixed synchronous/asynchronous designs.

On the verification front, we have verified several of our designs to be hazard-free using Burch's timed circuit verifier (Burch ICCD89) under the given timing constraints. Here, hazard-freedom is defined to mean that no transition once enabled to occur can be disabled without it occurring. This verification step serves as a double-check of our synthesis procedure.

We have also implemented the synthesis method in a CAD tool. This CAD tool begins with an event-rule system specification which is easily derived from many higher-level languages such as CSP, STG's, and the burst-mode specification. Timing analysis algorithms are then used to detect redundancies in the specification. Using the timing information, we then create a reduced state graph. From this reduced state graph, we then automatically derive a complex-gate implementation.

Future research for the next two quarters:

For the synthesis of speed-independent circuits, our current efforts are directed in completing the gate-level synthesis algorithm based on a realistic cell library. The current plan is to use the LSI logic cell library, which has been considered as an industry standard. Specifically, we must enhance the logic optimization procedures in the original synthesis program to allow for the breaking-up of high-fanin gates. Testability of speed-independent asynchronous circuits is also of primary interest for the next 6 months. Combining our previous effort in this direction [4,5], we believe that automatic synthesis of testable asynchronous circuits can be achieved at the gate-level.

For the synthesis of timed circuits, we plan to generalize our timing analysis algorithm to handle specifications with arbitrary non-deterministic behavior. Also, while we have verified our circuits to be hazard-free, in the future, we would like to verify that they meet the specifications as well. Other future goals are to synthesize several large examples and do the IC design of interesting timed circuits to better assess the area and performance gain compared with their synchronous counterparts.

References:

[1]. Peter A. Beerel and Teresa H.-Y. Meng, "Automated Gate-Level Synthesis of Speed-Independent Control Circuits", *1992 IEEE International Conference on Computer-Aided Design Digest of Technical Papers*, pp. 581-586, March 1992.

[2]. Chris Myers and Teresa H. Meng, "Synthesis of Timed Asynchronous Circuits", *Proc. of IEEE ICCD 1992*, pp. 279-284, October, 1992.

[3]. Chris Myers and Teresa H. Meng, "Synthesis of Timed Asynchronous Circuits", invited submission to *IEEE Trans. on VLSI Systems*, December 1992.

[4]. Peter A. Beerel and Teresa H.-Y. Meng. "Testability of Asynchronous Self-Timed Control Circuits with Delay Assumptions", *Proc. 28th ACM/IEEE Design Automation Conference*, pp. 446-451, June, 1991.

[5]. Peter A. Beerel and Teresa H.-Y. Meng, "Semi-Modularity and Testability of Speed-Independent Circuits", *Integration, the VLSI Journal*, Vol. 13, No. 3, pp.301-322, October 1992.